THE TRUE COST OF

AI-Driven Layoffs

Why They're Failing the Companies That Make Them

Chris Swecker
Director of Managed Services & AI
Appalachia Technologies
ChrisSwecker.com

May 2026

Executive Summary

Companies across every industry are reducing their workforces based on AI's speculative potential rather than its demonstrated performance. The data shows they are paying dearly for it — in rehiring costs, institutional knowledge loss, talent attrition, and, increasingly, stock price.

This whitepaper makes the case that committing to no AI-pretext layoffs is not an act of corporate charity. It is sound strategy. The evidence is consistent across independent research firms, academic surveys, and market data: organizations that cut staff in anticipation of AI capabilities that have not yet materialized are incurring costs that exceed their savings, often within six months.

The analysis draws on research from Harvard Business Review, Orgvue, Careerminds, Forrester, Gartner, McKinsey, Goldman Sachs, and Deutsche Bank. It also introduces the Cognitive Integration Spectrum — a five-level framework for classifying AI deployments by their actual depth of human-AI collaboration — as a practical decision tool for ensuring that workforce planning tracks deployment reality rather than deployment aspiration.

Key Findings

55%

of executives who made AI-driven layoffs already regret it

Orgvue, survey of 1,163 C-suite leaders, 2025

60% / 2%

cut headcount in anticipation of AI — only 2% tied cuts to actually deployed AI

HBR / Davenport & Srinivasan, survey of 1,006 global executives, Jan. 2026

42%

of companies have abandoned most AI initiatives entirely

S&P Global 451 Research, 2025 — up from 17% in 2024

−2%

average stock price movement following restructuring/automation layoff announcements, including AI-attributed cuts

Goldman Sachs analyst Elsie Peng, December 2025

6%

of enterprise leaders fully trust AI agents to autonomously run core business processes

HBR Analytic Services / Workato, survey of 603 business and technology leaders, 2025

Central Argument

Before any AI-driven workforce decision is made, leadership must ask a fundamental question: where is our AI actually operating on the Cognitive Integration Spectrum?

The Spectrum categorizes AI deployments across five levels of human-AI collaboration — from Level 1 (AI as Copyeditor) to Level 5 (AI as Proxy). The data shows that most organizations' AI implementations operate at Levels 1 through 3. Workforce decisions of any kind are only justified when AI is proven at scale at Level 4 or 5 — and even then, the IKEA model demonstrates that the better question is what higher-value work your people can do when AI handles the routine, not how many people you can eliminate.

1. The Deployment Gap: What Companies Thought AI Could Do

The foundational error behind most AI-driven layoffs is mistaking projection for performance. Companies did not eliminate roles because AI had proven itself capable of replacing human judgment at scale. They eliminated roles because they believed it would — soon.

A January 2026 Harvard Business Review survey of 1,006 global executives, conducted by Thomas Davenport of Babson College and MIT and Laks Srinivasan of the Return on AI Institute, delivered the most precise quantification of this gap yet: 60% of executives had already reduced headcount in anticipation of AI's future impact. Just 2% tied large layoffs to actual, deployed AI systems delivering results.

> *"When you ask CEOs who announced 'we're laying off 20% and replacing them with AI' whether they actually have a mature AI system in place — nine out of ten times the answer is no. They hadn't even started."*
>
> — J.P. Gownder, VP of Research, Forrester

The major research firms paint a consistent picture of where enterprise AI actually stands:

- Forrester's State of AI Survey (1,400+ global AI decision-makers): only 15% reported an EBITDA lift from AI in the past 12 months. Fewer than one-third could tie AI value to their organization's financial growth.
- McKinsey's State of AI Global Survey 2025 (1,993 participants): 88% of organizations use AI in at least one function, but approximately two-thirds remain in experiment or pilot mode. Only 39% attribute any measurable EBIT impact to AI. Just 6% qualify as "AI high performers."
- Gartner predicted that 30% of generative AI projects would be abandoned after proof of concept by end of 2025 — a figure multiple sources suggest proved conservative. By mid-2025, Gartner escalated, predicting over 40% of agentic AI projects would be canceled by 2027.
- S&P Global's 451 Research found 42% of companies had abandoned most AI initiatives in 2025, up from 17% the year prior. The average organization scrapped 46% of AI proofs of concept before reaching production.

IBM's CEO Arvind Krishna illustrated the phenomenon directly when he effectively acknowledged that much of what his company labeled "AI-driven cuts" was really post-pandemic over-hiring correction dressed in more palatable language. Deutsche Bank analysts, writing ahead of Davos 2026, gave the practice a name: "AI redundancy washing" — using the credibility of AI transformation to execute cost-cutting decisions that have little to do with deployed technology.

The ROI picture reinforces the deployment gap. Boston Consulting Group reports that 60% of firms see minimal gains from their AI investments. Deloitte's 2025 survey of over 1,800

executives found that while 85% are increasing AI investment, returns remain slow to materialize and hard to measure — with most executives reporting that satisfactory ROI on AI use cases takes two to four years. Cutting staff to fund or justify investments with a four-year payback horizon is not a transformation strategy. It is a gamble with someone else's livelihood.

The implication is significant: most organizations do not yet have the AI infrastructure to justify workforce restructuring decisions, let alone the institutional knowledge of how to deploy it effectively. The workforce decisions came before the deployments.

Gartner offers perhaps the starkest framing: according to their research, 0% of daily enterprise decisions were being made autonomously by AI as recently as 2024. Not a small percentage. Zero. Their forward projection — 15% of day-to-day work decisions made autonomously by AI by 2028 — is itself a signal that full autonomous deployment remains a future state, not a present one. Companies that cut staff in 2024 and 2025 based on autonomous AI capability were not early movers. They were operating on assumptions that the data shows were four years premature at minimum.

2. The True Cost Ledger

The financial case against AI-driven layoffs is not intuitive, because the savings are visible and immediate while the costs are diffuse and delayed. A payroll reduction shows up in the next quarterly report. The institutional knowledge that walked out the door, the survivor attrition it triggered, and the rehiring costs that followed show up months later — and rarely get attributed to the same decision.

When the full ledger is examined, the picture changes considerably.

Regret Is the Majority Position

Orgvue's April 2025 survey of 1,163 C-suite and senior decision-makers across the US, UK, Canada, Australia, and Asia found that 39% of business leaders made employees redundant as a result of deploying AI. Of those, 55% admit they made wrong decisions about those redundancies. This is not a fringe outcome. It is the statistical majority of executives who made AI-driven workforce cuts, already wishing they had done otherwise.

Additional data from the same survey: 34% of leaders admit employees quit as a direct result of AI introduction. 25% acknowledge they don't know which roles in their organization can benefit most from AI. Orgvue CEO Oliver Shaw stated plainly: "Businesses are learning the hard way that replacing people with AI without fully understanding the impact on their workforce can go badly wrong."

The Rehiring Numbers

Careerminds surveyed 600 HR professionals in February 2026 who had conducted AI-led layoffs in the prior twelve months. The findings on financial outcomes are stark:

- 35.6% had already rehired for more than half of the roles they eliminated.
- 52.1% rehired within six months of the original cuts.
- 30.9% found bringing roles back cost more than they saved by cutting them.
- 42.4% broke even — meaning the rehiring entirely canceled out the savings.
- Only 26.7% came out ahead financially.
- Only 8.4% said the AI-driven restructuring delivered what was promised.

In other words, roughly three-quarters of organizations that made AI-driven cuts ended up financially neutral or worse once the rehiring cycle completed — and they incurred all the operational disruption, morale damage, and knowledge loss in between.

The Institutional Knowledge Problem — and the AI Degradation Loop

The most consistently underweighted cost of layoffs — AI-driven or otherwise — is the institutional knowledge that leaves with the people. A Panopto/YouGov study of 1,001 U.S. workers found that 42% of institutional knowledge is unique to the individual employee and not shared with colleagues. Knowledge workers waste an average of 5.3 hours per week waiting for or recreating knowledge that should already exist. For a large organization, this translates to an estimated $47 million annually in productivity loss from inefficient knowledge sharing alone.

What this means in the AI context: no matter how capable the AI system, it can only process what has been documented and fed to it. It cannot recover what existed only in the minds of people who have been let go. The Careerminds data confirms this: 32.9% of companies that made AI-driven cuts reported losing critical skills and expertise when those employees left. Only 21.4% said AI fully replaced the eliminated roles without operational issues.

But the damage runs deeper than a one-time knowledge loss. There is a compounding effect that almost never appears in layoff post-mortems: when the people who held edge-case knowledge leave, the AI system itself begins to degrade. The remaining staff — overextended, demoralized, and now operating without the colleagues who understood the nuances — are the ones prompting, training, and correcting the AI. Their inputs become the ground truth. And if they don't know what they don't know — because the person who knew left — the AI learns from increasingly incomplete context. The model doesn't fail dramatically. It drifts. Outputs become subtly worse. Edge cases get mishandled. The system that was supposed to replace human judgment is now being degraded by the absence of it. You didn't just lose the employee. You poisoned the AI they were meant to hand off to.

Note: This degradation loop is analytical inference drawn from the Panopto/YouGov and Careerminds data cited above; it has not, to the author's knowledge, been formally studied as a discrete phenomenon.

Restructure Contagion: The Survivor Attrition Multiplier

Researchers have begun using the term "restructure contagion" to describe a secondary wave of voluntary departures that compounds the damage of the original cuts — one of the most expensive and least-discussed consequences of AI-motivated workforce reductions, because it is delayed, diffuse, and rarely attributed to the decision that caused it.

Academic research on layoffs consistently demonstrates that the financial damage extends well beyond the eliminated roles. Trevor and Nyberg's landmark 2008 study, published in the Academy of Management Journal, found that layoffs targeting just 1% of a workforce led to a 31% increase in voluntary turnover among surviving employees. In organizations with low procedural justice, the predicted increase reached 112%.

The mechanism is straightforward: high-performing employees — the ones with the most options, and the ones organizations can least afford to lose — rationally update their calculation after watching colleagues lose their jobs. They do not wait to see if they are next. They begin looking.

A 2024 HBR/Culture Amp study of 146 companies found that after layoffs, company confidence dropped 16.9 points, belief in career opportunity dropped 12.1 points, and it takes 12–24 months for engagement to recover. A Glassdoor study estimated $20.8 billion in collective losses in the first year following layoffs across 197 companies, with a 40% jump in current employees actively searching for new roles. The Careerminds survey found 47% of respondents reported productivity declines after cuts, and 41% saw additional talent leave through resignation. Deloitte research found that losing just eight key employees voluntarily following a ten-person

layoff can cost a company $1.2 million in unplanned expenses — making the restructure more expensive than doing nothing.

None of these costs appear in the original layoff ROI calculation. All of them are real.

The True Replacement Cost: A Five-Phase Accounting

Even without the AI-specific complications, the raw cost of replacing employees is substantial — and consistently underestimated because the standard calculation only captures the most visible phase. SHRM's methodology identifies five distinct cost phases in a full layoff-and-rehire cycle:

- Separation Costs: Severance packages (averaging 38% of salary when benefits are included), unemployment insurance contributions, and administrative exit processing.
- Vacancy Costs: Lost revenue and productivity while the role sits empty. For revenue-generating positions, SHRM estimates this at up to triple the daily compensation rate.
- Recruitment Expenses: External agency fees (typically 20% of annual salary), advertising, and the significant internal labor hours of hiring managers and recruiters.
- Onboarding and Training: The productivity loss of peer trainers who must sacrifice their own output to bring the new hire up to speed, plus equipment and systems costs.
- Productivity Ramp-Up: New employees take six months to two years to reach the productivity level of the person they replaced. This cost is real and almost never appears in a layoff ROI model.

Applied to a $100,000 position, SHRM's framework produces the following conservative estimate:

Cost Phase	Calculation Method	Estimated Cost
Separation	Severance (2 wks) + 38% benefits load + admin	~$12,500
Vacancy (30 days)	Daily rate ($400) × 30	~$12,000
Recruitment	Agency fee (20%) + ads + recruiter hours	~$25,000
Onboarding	Trainer time + equipment + licenses	~$7,500
Ramp-Up (6 months)	50% productivity loss × 125 working days	~$25,000
Total Cycle Cost	**Sum of all phases**	**~$82,000**

That $82,000 figure is conservative and covers a single mid-level position. Multiplied across a restructuring event of 100 people with a 35.6% rehire rate, the unplanned replacement cost reaches approximately $2.9 million — often eclipsing the projected salary savings entirely, before survivor attrition costs are factored in.

Wharton School research adds another dimension: external hires command 18–20% higher salaries than internal candidates and underperform comparable internal hires for their first two years. McKinsey found that in approximately 75% of cases, it is cheaper to reskill an existing employee than to hire externally — and that effective reskilling delivers a 6–12% productivity

uplift. The math, taken whole, strongly favors investment in existing people over reduction and replacement.

3. What the Market Is Telling You

For years, announcing layoffs — even when framed as AI-driven transformation — was rewarded by markets as evidence of decisive leadership and cost discipline. That dynamic has reversed. The investor community has moved through three distinct phases on this question: rewarding the announcement, then growing skeptical of the claim, and now actively pricing in the risk of premature AI-attributed cuts.

The Signal Has Inverted

Goldman Sachs analyst Elsie Peng published research in December 2025 analyzing post-announcement stock performance for companies citing automation and restructuring as layoff drivers — a category that includes, but is not limited to, AI-attributed cuts. The finding: stocks fell by an average of 2% following these announcements, a direct inversion of the historical pattern in which restructuring signals were treated as bullish. The Goldman analysis identified why: firms making these announcements showed higher capex, higher debt, higher interest expense, and lower profit growth than comparable companies. Investors have concluded that AI-attributed layoff language is more likely to signal financial distress dressed as innovation than genuine AI-driven efficiency.

The Practice Has a Name

Deutsche Bank Research, writing ahead of the January 2026 World Economic Forum, formalized the investor concern with a term: "AI redundancy washing." Their analysts defined it as the practice of using AI transformation as cover for cost-cutting decisions that have little relationship to actual technology deployment, and predicted it would be a defining feature of 2026. The implication for investors: treat AI layoff announcements with significant skepticism until production deployment evidence is provided. The J.P. Morgan Asset Management data gives that skepticism empirical grounding — AI has been cited in approximately 55,000 of 1.1 million job cuts announced in 2025, less than 5% of the total. The narrative vastly outpaces the reality.

Capital Is Repricing Toward Retention

The investor shift is not only about skepticism toward layoffs. It is also an active repricing toward organizations that invest in their people alongside AI. Mercer's Global Talent Trends 2026 survey of approximately 12,000 executives, HR leaders, investors, and employees found that 97% of investors said their funding decisions would be negatively impacted by organizations failing to adopt agile talent models. 77% were more likely to invest in companies committed to empowering employees through AI education. The signal from the investment community has shifted from "how lean can you run" to "how effectively are you combining humans and AI — and can you prove it?"

Gartner's operational data anchors the Mercer investor sentiment in business reality: by 2027, Gartner predicts half the companies that cut customer service staff citing AI will have to rehire. Their survey of 321 customer service leaders found only 20% had actually reduced staffing because of deployed AI. The rest were anticipating — the same projection-vs-performance gap that the HBR survey documented in the C-suite.

What the investor community is sensing, the enterprise data confirms: trust in autonomous AI is not growing with adoption — it is declining. Capgemini's 2025 survey of 1,500 senior leaders found that trust in fully autonomous AI agents dropped from 43% to 27% in a single year, with 71% of organizations stating they cannot fully trust autonomous AI agents for enterprise use. The HBR Analytic Services/Workato survey of 603 business and technology leaders found that only 6% fully trust AI agents to autonomously run core business processes. Another 43% trust agents only with limited or routine operational tasks. The lived experience of deploying AI in production is making organizations more cautious, not less — which is precisely the opposite of the story being told in layoff announcements.

The Long-Term Financial Record

The academic longitudinal data completes the picture. A 2024 meta-analysis published in Frontiers in Behavioral Economics, encompassing 905 effect sizes across decades of downsizing research, found no significant positive relationship between workforce reductions and market performance even two years after the event. Lindenwood University longitudinal studies found that while a layoff announcement can temporarily inflate a stock's price-to-earnings ratio by signaling future cost reductions, actual return on equity does not improve and often enters a steeper downward trend two years later. Most tellingly, measures of human capital productivity — specifically profits per employee and sales per employee — show a significant downward trend in companies that announce layoffs compared to those that do not.

Taken together, the evidence builds a single argument: the financial market has caught up with the operational reality documented in Sections 1 and 2. AI-attributed layoffs are not rewarded as bold transformation moves. They are increasingly read as a signal that the underlying business is in trouble, and that management is using AI as cover. The short-term pain for long-term gain narrative is not supported by the longitudinal evidence. The pain is real. The long-term gain largely is not.

> *"AI redundancy washing will be a significant feature of 2026."*
>
> — Deutsche Bank Research Institute, January 2026

4. What It Looks Like When You Get It Right

The argument against AI-driven layoffs is stronger when it comes with an alternative model. Three organizations — at different scales and in different industries — demonstrate what responsible AI workforce transition can produce.

IKEA: The Cost Center That Became a Profit Center

In 2021, IKEA's franchisee Ingka Group deployed an AI chatbot called Billie — named after their iconic Billy bookcase — to handle routine customer service inquiries. By 2023, Billie was resolving approximately 47% of customer queries, representing 3.2 million interactions and nearly €13 million in savings annually. By most organizational playbooks, this is the point at which you announce headcount reductions and take credit for AI efficiency.

IKEA asked a different question: what could these people do if they were no longer handling routine inquiries?

The answer was interior design consulting. Ingka retrained 8,500 call center workers as remote design advisors, building on their existing customer knowledge and product familiarity and adding structured design training. The reskilled workers now staff a remote consulting channel that generates approximately €1.3 billion (~$1.4B) in annual sales — 3.3% of total revenue, with a stated goal of 10% by 2028. The company did not reduce headcount. It eliminated a cost center and created a revenue line.

> *“The chatbot handles "where's my order?" The humans, armed with design training and AI-powered tools, help customers reimagine their living spaces.”*

Ericsson: Reskilling at Enterprise Scale

Ericsson reskilled more than 15,000 employees in AI and data science over three years, beginning around 2018, transforming telecommunications experts into AI and data science specialists. The program was documented in a Harvard Business Review study by Harvard's Digital Reskilling Lab and BCG Henderson Institute across approximately 40 organizations. A distinguishing feature of the Ericsson approach: reskilling was embedded in OKRs and reviewed quarterly by executives, treating workforce transformation as an operational priority rather than an HR initiative.

JPMorgan Chase: Investment as Strategy

JPMorgan Chase CEO Jamie Dimon stated at a February 2025 investor meeting: "We already have huge redeployment plans for our own people... we have displaced people from AI, and we offer them other jobs." The bank invests $300 million per year in employee training, with an additional $350 million in global workforce investment. Some 150,000 employees now use the bank's internal LLM platform weekly. A controls review process formerly requiring 200 people

was automated — and the bank identified 3,000–5,000 additional employees who could benefit from similar tools, with redeployment rather than elimination as the stated outcome.

A Mid-Market Blueprint

The honest acknowledgment first: the documented case studies in this space are dominated by Fortune 500 organizations with substantial training budgets. Verified, outcome-rich examples from mid-market companies ($100M–$1B revenue) are limited in the public record. KeyBank's partnership with Tech Elevator — which reskilled 80+ employees with 100% retention — suggests the model holds at smaller scale, but the evidence base is thin. This is a genuine gap in the research, and smaller organizations should build their own data points rather than assume the macro evidence translates directly.

What does translate is the structural logic. The IKEA case did not succeed because IKEA is large. It succeeded because leadership asked the right question before defaulting to cuts: what could these people do if AI handled the routine work? That question is free. It costs nothing to ask, and the answer is frequently more valuable than the salary you were considering eliminating.

For organizations that cannot afford enterprise-scale reskilling programs, here is a practical starting framework:

- Designate one internal champion per affected team. Their job: map current AI deployments against the Cognitive Integration Spectrum and identify adjacent higher-value work the team could take on. This requires time, not budget.
- Apply the 5:1 principle at your scale. MIT research found that firms spending $5 on people for every $1 on technology are four times more likely to maintain top-tier financial performance. At smaller scale, that might mean pairing each AI tool rollout with a structured 30-day internal learning sprint — not a formal training program, but a deliberate period where team members experiment, document, and share what they are learning. The ratio matters more than the dollar amount.
- Pilot before you restructure. Before eliminating any role, run a 90-day redeployment pilot: give the affected employee a defined set of higher-value responsibilities, AI tools to support them, and a clear success metric. The cost of the pilot is 90 days of existing salary. The cost of getting it wrong is the full SHRM replacement cycle plus everything that leaves with them.
- Document as you go. The mid-market organization's greatest disadvantage is that institutional knowledge is even more concentrated in fewer people. Before any AI implementation, invest in knowledge capture — process documentation, decision logs, edge-case libraries. This protects the AI deployment and creates the reskilling foundation at the same time.

The investment levels differ from what IKEA or JPMorgan spent. The principles do not.

5. The Cognitive Integration Spectrum: A Decision Framework

The central failure behind most AI-driven layoffs is a category error: treating AI capability projections as if they were deployment realities. The Cognitive Integration Spectrum provides a structured way to prevent that error from entering a workforce planning conversation.

The Spectrum classifies AI use cases by the actual depth of human-AI collaboration in a given deployment — not by what the technology could theoretically do, but by what it is actually doing in production today. It runs from five levels:

Level	Action / Role	Description	Examples	Workforce Decision Risk
Level 1	**Enhance** AI as Copyeditor	AI polishes or formats human output. Human does all thinking; AI improves presentation.	Email optimization, writing polish, tone adjustment	Zero
Level 2	**Extend** AI as Intern	AI handles lookup, brainstorming, or legwork. Still needs direction; human decides what to do with the result.	Research, troubleshooting partner, transcription	Zero
Level 3	**Collaborate** AI as Colleague	Iterative back-and-forth where both parties shape the output.	Documentation assistant, use case ideation, review drafting	Low — role evolution warranted
Level 4	**Delegate** AI as Direct Report	Human defines the task, AI executes end-to-end, human reviews the result.	Meeting notes, ticket triage, knowledge base Q&A	Moderate — only if proven at scale
Level 5	**Orchestrate** AI as Proxy / Deputy	Multi-step autonomous workflows with checkpoint-based human review, not step-by-step oversight.	Agent workflows, vCAIO planning, automated onboarding	Warrants workforce planning — rare in practice today

The Data Behind the Spectrum

The Cognitive Integration Spectrum was developed as a practical decision tool, but it turns out to be an empirically accurate one. Independent research from multiple major firms has arrived at nearly identical frameworks — which means the CIS is not just a useful model, it maps to documented deployment reality.

BCG's 4-Tier Autonomy Framework, published in 2025, is structurally parallel to the CIS: Shadow Mode (AI runs in parallel, humans decide) — Copilot (AI assists, humans approve) — Supervised Autonomy (AI acts, humans oversee) — Full Autonomy (AI operates independently). BCG states explicitly: "Right now, the vast majority of companies are building most processes with humans in the loop and with ultimate decision authority." Full autonomy is "reserved for

highly mature, low-risk environments where the cost of error is negligible." Two independent frameworks, developed separately, converging on the same five-level structure and the same conclusion about where most deployments actually sit.

The empirical distribution data validates both. Across 15+ major surveys from 2024–2026 — Gartner, McKinsey, Deloitte, BCG, Capgemini, IDC, HBR, and others — the consensus is unambiguous: 85–94% of enterprise AI deployments operate with meaningful human oversight. Fully autonomous AI execution remains confined to 2–7% of organizations depending on definition stringency. The specific data points:

- Capgemini (1,500 senior leaders, 2025): only 2% have deployed AI agents at full scale. 61% are still in exploration. Trust in fully autonomous agents dropped from 43% to 27% in a single year.
- IDC/AWS (900+ enterprises, 2025): 79% have adopted AI agents in some form, but only 2.9% are scaling agentic use cases across departments.
- HBR Analytic Services/Workato (603 leaders, 2025): only 6% fully trust AI to run core business processes autonomously. 39% restrict AI to supervised or non-core use cases.
- Gartner (2025): fewer than 5% of enterprise applications have task-specific AI agents today, with the vast majority still operating at the assistant stage.
- MIT CISR (721 companies, 2024 baseline; 2025 update): the 2024 baseline showed 62% in Stages 1–2. The 2025 update shows meaningful movement — Stage 3 (developing AI ways of working) grew from 31% to 46%, and Stage 4 (AI future-ready) grew from 7% to 18%. This is real progress, but it does not contradict the Levels 1–3 dominance finding: MIT explicitly notes that Stage 3 means industrializing workflows and scaling AI — not autonomous execution. Even Stage 4 organizations maintain structured human oversight for mission-critical decisions.

This data maps directly onto the Cognitive Integration Spectrum. The 85–94% of organizations operating with human oversight are, by definition, operating at Levels 1 through 3 — Enhance, Extend, and Collaborate. The 2–7% with genuine autonomous deployment are approaching Levels 4 and 5. The workforce decisions being made based on Level 4–5 assumptions are being made in an enterprise environment that is, overwhelmingly, still operating at Level 1 through 3.

That is not a temporary gap that will close quickly. Gartner's projection that 15% of daily work decisions will be made autonomously by 2028 — up from 0% in 2024 — is an optimistic forecast from a firm with strong incentives to project AI adoption. Even if accurate, it means that four years from now, 85% of daily decisions will still involve human judgment. The workforce decisions being made today will outlast the AI deployment assumptions they were based on.

Why Five Levels, Not Four

The BCG 4-Tier Autonomy Framework is the closest published analog to the Cognitive Integration Spectrum, and the convergence between them is meaningful — two frameworks developed independently arriving at the same basic structure and the same conclusions about where most enterprise AI actually sits. But the structural difference is worth examining, because it reflects a difference in purpose.

BCG’s framework was designed to describe the technical architecture of an AI deployment: how much autonomy the system has, and what kind of human oversight is in place. For that purpose,

four tiers are sufficient. Their Copilot tier — AI assists, humans approve — captures a broad range of deployment types under a single label.

The Cognitive Integration Spectrum was designed for a different purpose: governing the workforce planning conversation. And for that purpose, the middle of the range requires more precision than BCG's framework provides.

What BCG calls Copilot encompasses two meaningfully different relationships between humans and AI. At Level 2 (Extend), the AI is handling discrete tasks — research, lookup, transcription — on human direction, and the human decides entirely what to do with the result. The AI is a capable assistant, but the thinking is one-directional. At Level 3 (Collaborate), something different is happening: the human and AI are in iterative back-and-forth, each shaping the other's output, with the final product emerging from that exchange rather than from the human alone.

Level 3 is also the point at which the human-AI contribution ratio begins to approach parity — the first level where an outside observer might reasonably ask how much of the output the human actually produced. This makes it a genuine tipping point in the deployment arc. But the tipping point runs in the opposite direction than it appears. What leadership often reads as the employee becoming redundant is actually the employee becoming more productive — their judgment, context, and domain knowledge are being amplified by the AI, not replaced by it. The output looks different because the ceiling has been raised, not because the human has been removed from the equation.

That misread is the most common source of premature workforce decisions. Level 2 warrants no restructuring discussion whatsoever. Level 3 warrants a role evolution conversation: how is this job changing, what new skills does it require, and how do we invest in the people doing it? Collapsing those two into a single tier — as BCG's framework does — obscures exactly the question that matters most at the moment it matters most.

The top two tiers align closely. BCG's Supervised Autonomy corresponds to Level 4 (Delegate), and Full Autonomy corresponds to Level 5 (Orchestrate). Both frameworks agree that genuine autonomous execution is rare in production environments today, and both reserve their highest tier for systems operating independently with only checkpoint-based human review.

The CIS is not a refinement of BCG's framework — it was developed separately for a different use case. But the convergence at the extremes and the deliberate divergence in the middle reflect the same underlying reality: the action, for most organizations, is happening in the collaborative range. That is where the category errors get made, and where a more precise vocabulary pays off.

How to Use the Spectrum in Workforce Decisions

The diagnostic question is simple: before any AI-driven workforce decision is brought to leadership, require the presenting team to map the relevant AI deployment on the Spectrum. Then apply the following logic:

- Level 1–2 deployments (Enhance / Extend): AI is assisting human output or handling discrete lookups. Zero workforce restructuring is justified at these levels. The AI does not replace human judgment; it supports it.
- Level 3 deployments (Collaborate): AI and humans are iteratively co-producing work. Role evolution conversations are warranted — discussions about how the job is changing,

what new skills are needed, and where human contribution is being amplified. Not headcount reduction.

- Level 4–5 deployments (Delegate / Orchestrate): AI is executing substantive workflows end-to-end with human review. These deployments are rare in genuine production settings today and are the only level at which workforce planning is legitimately on the table — and even then, the IKEA question should precede any reduction: what higher-value work can these people do that AI cannot?

A Worked Example

Consider a finance team that has deployed an AI tool to assist with accounts payable review. The tool flags anomalies, suggests categorizations, and drafts exception reports. A human reviews every output and makes every final call. That is a Level 3 deployment — Collaborate. AI and human are iteratively shaping the work together.

The appropriate workforce conversation at Level 3 is not "how many AP clerks can we eliminate?" It is: "Now that our AP staff spends less time on routine categorization, what higher-value financial analysis work are they available to do? What would it take to develop their skills to operate at that level? And what does the AP function look like in twelve months if we invest in that path?"

If the same team then builds a fully automated AP workflow that processes invoices end-to-end with weekly human audit reviews — and that system has been running reliably in production for six months with documented error rates — that is approaching Level 4. Now a workforce planning conversation is legitimate. But the first question is still the IKEA question: what can these people do that the system cannot? Vendor relationship management, exception escalation, financial strategy, client-facing work. The answer is usually more than the spreadsheet assumes.

The data suggests that most organizations' AI deployments currently operate at Levels 1 through 3. Forrester found that only 16% of individual workers had high AI readiness (AIQ) in 2025. McKinsey found two-thirds of organizations remain in experiment or pilot mode. Gartner found that 42% had abandoned most AI initiatives before reaching production. The workforce decisions being made today are being made against a Level 1–2 deployment reality while imagining a Level 4–5 future.

There is a deeper reason why existing employees are the right people at Levels 3 through 5 — and it has nothing to do with sentiment. Someone inside your organization who has developed AI fluency is more impactful than an AI expert who knows nothing about your business. The higher the level of integration, the more the human's role is to supply judgment, context, and institutional knowledge that the AI cannot generate on its own. That requires knowing the company, the clients, the edge cases, and the history. An outside AI specialist, however technically capable, arrives without any of it. Your existing people already have it. Training them is not charity — it is the more effective path to AI deployment that actually works.

BCG research provides the empirical grounding for why human judgment remains irreplaceable at the higher levels. Their 2025 analysis found that AI generates substantial productivity improvements in structured, bounded tasks — 14 to 40%, and up to 96% in certain writing tasks. But for complex, creative, or strategic work, those productivity gains largely vanish. This is precisely why deployment so frequently disappoints at scale: organizations attempt to apply assistive or early-stage agentic tools to problems that require the kind of contextual reasoning that only comes from experience inside the organization. MIT’s NANDA initiative quantified this directly in their “The GenAI Divide: State of AI in Business 2025” report — analyzing 300 AI

deployments across 150 companies, they found that 95% of AI pilot projects failed to deliver measurable P&L impact. The failure was rarely the AI model itself; it was the organizational mismatch between tool capability and task complexity. The Cognitive Integration Spectrum exists to prevent that same mismatch from appearing in workforce planning decisions.

> *"Most companies are making Level 5 workforce decisions based on Level 2 AI deployments. The Spectrum makes that gap impossible to ignore."*

The Spectrum as an Accountability Tool

Beyond individual decisions, the Spectrum serves as an ongoing accountability mechanism. When each AI use case in the organization is mapped against its actual operating level, the gap between where leadership believes the organization is and where it actually is becomes visible. That visibility is the precondition for honest workforce planning, honest AI investment decisions, and honest communication with employees, investors, and clients.

The goal is not to prevent progress toward higher levels of integration. The goal is to ensure that workforce decisions track deployment reality rather than deployment aspiration.

6. Recommendations

The following five recommendations are drawn directly from the evidence in this whitepaper. They are not aspirational. They are the organizational behaviors that distinguish the companies succeeding with AI from the ones navigating expensive reversals.

01

Map before you cut

Before any workforce decision, require AI leads to map every AI deployment on the Cognitive Integration Spectrum. If the deployment is operating at Level 1 or 2, the workforce case does not exist yet. End of discussion. This mapping discipline also addresses a growing governance problem: a 2025 AuditBoard and Panterra survey of 400+ GRC professionals found that 82% are using AI across business functions, yet only 25% have a fully implemented AI governance program. Organizations are running autonomous systems against real data without clear accountability structures. Mapping deployments to the Spectrum is a practical first step toward closing that gap.

02

Demand production evidence

Ask these questions before approving any AI-attributed headcount reduction: What AI system is deployed in production and proven at scale? What are its error rates? What is the rollback plan? If the answers involve pilots, future plans, or vague efficiency references, the layoff is premature.

03

Ask the IKEA question

Before eliminating a role, ask: what could this person do if we freed them from the routine work? Customer service workers became $1.4B-generating interior design consultants. The right question is not "can AI do this job?" but "what higher-value work does this person become capable of?"

04

Require a reskilling plan before cuts are approved

Per the Careerminds data, 55% of organizations that made AI-driven cuts never formally discussed reskilling as an alternative. Make it a required step in the decision process, not an afterthought. MIT research found that firms spending $5 on people for every $1 spent on technology are four times more likely to maintain top-tier financial performance. And McKinsey's 2025 survey found that 65% of AI high performers — organizations generating more than 5% of EBIT from AI — have defined formal human-in-the-loop validation processes, compared to just 23% of other organizations. Maintaining human oversight is not a sign of AI immaturity. It is what the most successful AI deployments actually do.

05

Make the commitment public

A no-AI-pretext-layoff pledge is a talent acquisition strategy, a retention signal, and a market accountability mechanism. Make it explicit and visible. It forces honest internal accounting, signals to prospective employees that growth rather than fear defines your relationship with technology, and tells investors that your AI strategy is built on deployment reality rather than speculation. For leaders not yet ready to make a public pledge: start internally. Commit to your current team first — in writing, in a town hall, in the words your managers use. The Trevor & Nyberg research shows that procedural justice — how decisions are made and communicated — is as important as the decisions themselves in preventing voluntary turnover. An internal commitment builds that trust and creates the accountability structure that makes a future public pledge credible rather than performative.

NOTE: IF REDUCTIONS OCCUR FOR OTHER LEGITIMATE REASONS

If workforce reductions do occur for genuine structural or business reasons — not AI pretext, but real operational change — treat survivor health as a measurable output, not an afterthought. Monitor voluntary turnover rates, engagement scores, and productivity among remaining staff for a full 12 months post-reduction. Trevor & Nyberg's research shows a 31% increase in voluntary turnover follows even a 1% workforce reduction. That number almost never appears in the original layoff ROI model. It should be a required line item in every post-reduction review.

Conclusion

The argument in this paper is not a call for moral courage. It is a call for accurate accounting.

The companies that made aggressive AI-driven workforce reductions in 2023 and 2024 were operating on a set of assumptions — about what AI could do, how quickly it could be deployed, and how the market would respond — that have not held up. The research is now clear: the majority regret it, most broke even or lost money on the rehiring cycle, the best people started leaving, and the market has reversed its reward signal for the behavior.

The companies that got it right asked a different question. Not "how many roles can AI eliminate?" but "what becomes possible for our people when AI handles the routine?" IKEA turned customer service into a $1.4 billion revenue channel. Ericsson turned telecom engineers into AI and data scientists. JPMorgan redeployed employees rather than eliminate them and built 150,000 AI-enabled workers in the process.

The Cognitive Integration Spectrum gives leadership a language and a framework for having this conversation honestly. It makes the gap between projection and deployment visible. It prevents the category error of making Level 5 decisions based on Level 2 realities.

And it points toward the right person for the seat next to the AI. Not an outside expert imported for their technical credentials. Someone inside your organization who knows your clients, your processes, your history, and your edge cases — and who has learned to work with AI on top of that foundation. A person with deep domain knowledge and developing AI fluency will outperform an AI specialist who knows nothing about your business every time — as Mark Leeper, Senior Manager of AI Adoption at Hatz.AI, has observed. The institutional knowledge that makes someone dangerous with AI is exactly the same institutional knowledge that walks out the door when you cut them.

A no-AI-pretext layoff commitment is not a values statement. It is a strategy. It is what the evidence supports. And for what it is worth, it also happens to be the right thing to do.

Those two things are, in this case, the same thing.

Sources & Verification Notes

The following sources are the primary basis for factual claims in this whitepaper. Sources marked (a) were confirmed via direct access to the original publication. Sources marked (b) were confirmed via consistent secondary reporting across multiple credible outlets; original research is behind institutional paywalls.

Core Survey Data

- Orgvue / Vitreous World (2025). "Human-first, machine enhanced: From optimism to pragmatism in AI-driven workforce transformation." Surveyed 1,163 C-suite and senior leaders, Feb–Mar 2025. orgvue.com/news/55-of-businesses-admit-wrong-decisions [Confirmed at source — (a)]
- Davenport, T.H. & Srinivasan, L. (2026, January 29). "Companies Are Laying Off Workers Because of AI's Potential — Not Its Performance." Harvard Business Review. hbr.org/2026/01/companies-are-laying-off-workers-because-of-ais-potential-not-its-performance [Article existence confirmed at source; statistics confirmed via secondary reporting — (a)/(b)]
- Careerminds (2026, February 19). "AI-led layoffs: What HR leaders wish they knew before making job cuts." Survey of 600 HR leaders, fielded Feb 12–14, 2026. careerminds.com/blog/cost-of-ai-layoffs [Confirmed at source — (a)] Note: Careerminds is an outplacement services provider with commercial interest in highlighting layoff costs.

Analyst Research

- Forrester Research. "Predictions 2026: The Future of Work." Including AIQ data and 9-in-10 CEO finding. forrester.com [Confirmed via Forrester publications and secondary reporting — (a)/(b)]
- Gartner (2024–2025). Various press releases on AI project abandonment rates, generative AI predictions, and 5-Stage Agentic AI Evolution timeline. gartner.com/en/newsroom/press-releases [Confirmed at source — (a)] Note: 0% autonomous decision-making baseline (2024) and 15% projection (2028) from Gartner research.
- McKinsey & Company (2025). "The State of AI: Global Survey 2025." 1,993 participants, 105 nations. mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai [Confirmed at source — (a)] Includes finding that 65% of AI high performers maintain formal human-in-the-loop validation vs. 23% of others.
- S&P Global 451 Research (2025). Voice of the Enterprise AI survey, 1,006 respondents. spglobal.com/market-intelligence [Confirmed at source — (a)]
- Deloitte (2025). State of AI in the Enterprise 2026. n=3,235 senior leaders, 24 countries, Aug–Sep 2025. deloitte.com [Confirmed via secondary reporting — (b)]
- BCG (2025). AI at Work 2025 and 4-Tier Autonomy Framework (Shadow → Co-pilot → Supervised → Full Autonomy). Includes productivity paradox finding: AI gains of 14–40% in structured tasks; gains largely vanish for complex/strategic work. bcg.com [Confirmed via secondary reporting — (b)]
- MIT NANDA Initiative (2025). "The GenAI Divide: State of AI in Business 2025." Analysis of 300 AI deployments across 150 companies; 95% of AI pilot projects delivered no measurable P&L impact. Reported by Fortune (fortune.com/2025/08/18/mit-report-95-percent-generative-ai-pilots-at-companies-failing-cfo/). [Confirmed via multiple outlets — (b)] Note: methodology uses a narrow definition of success (measurable P&L impact within 6 months); broader ROI measures would yield different results.
- AuditBoard / Panterra (2025). "From blueprint to reality: Execute effective AI governance in a volatile landscape." Survey of 400+ GRC professionals. Key finding:

82% using AI across functions; only 25% have fully implemented governance programs. auditboard.com/blog/benchmarking-ai-governance-4-key-survey-findings [Confirmed at source — (a)]

- Capgemini (2025). Rise of Agentic AI report. n=1,500 senior leaders, 14 countries, July 2025. capgemini.com [Confirmed via secondary reporting — (b)] Key finding: trust in fully autonomous agents dropped from 43% to 27% over one year; only 2% have deployed AI agents at full scale.
- IDC / AWS (2025). Enterprise AI agent adoption study. n=900+ enterprises, November 2025. idc.com [Confirmed via secondary reporting — (b)]
- HBR Analytic Services / Workato (2025). Survey of 603 business and technology leaders on AI agent trust, July 2025. hbr.org [Confirmed via secondary reporting — (b)] Key finding: only 6% fully trust AI agents to autonomously run core business processes.
- MIT CISR (2024/2025). Enterprise AI Maturity Model. n=721 companies, 2024 baseline with 2025 update. cisr.mit.edu [Confirmed via secondary reporting — (b)] Note: 2025 update shows Stage 3 grew to 46% and Stage 4 to 18%; MIT confirms higher stages do not equate to autonomous execution.

Market & Investor Data

- Goldman Sachs, Peng, E. (December 2025). Research note on post-announcement stock performance following AI-attributed layoffs. Reported by Fortune (fortune.com/2025/12/25/goldman-sachs-research-ceos-layoffs-stock-price/). [Behind institutional paywall; confirmed via secondary reporting across Fortune, Yahoo Finance, Benzinga — (b)] Note: ~2% figure applies to restructuring/automation layoffs broadly, not exclusively AI-only layoffs.
- Deutsche Bank Research Institute, Cox, A. & Abrudan (January 20, 2026). Published ahead of World Economic Forum, Davos. "AI redundancy washing" quote confirmed via CNBC reporting. [Behind client paywall; confirmed via secondary reporting — (b)]
- J.P. Morgan Asset Management (2025). Analysis citing AI attribution in approximately 55,000 of 1.1 million job cuts announced in 2025. am.jpmorgan.com [Confirmed at source — (a)]
- Mercer (2026). Global Talent Trends 2026. Survey of ~12,000 executives, HR leaders, investors, and employees. mercer.com [Confirmed at source — (a)]

Case Studies

- IKEA / Ingka Group (2023, June 29). Official newsroom press release on Billie chatbot and reskilling program. Confirmed via Ingka Group Newsroom and Reuters (Helen Reid, June 13, 2023). Note: €1.3B (~$1.4B) represents total sales through the remote consulting channel, not net incremental revenue solely attributable to reskilling. [Confirmed at source — (a)]
- Ericsson reskilling program. Documented in: Hidalgo, C. et al. (2023, September). "Reskilling in the Age of AI." Harvard Business Review. hbr.org/2023/09/reskilling-in-the-age-of-ai [Confirmed via HBR secondary reporting — (b)]

- JPMorgan Chase. CEO Jamie Dimon, investor meeting statements, February 2025. Corporate training investment figures confirmed via press releases and corporate communications. [Confirmed via public statements — (b)]

Academic Research

- Trevor, C.O. & Nyberg, A.J. (2008). "Keeping Your Headcount When All About You Are Losing Theirs." Academy of Management Journal, 51(2), 259–276. [Peer-reviewed; confirmed at source — (a)]
- Panopto / YouGov (2018). Workplace Knowledge and Productivity Report. 1,001 U.S. workers. panopto.com/resource/ebook/valuing-workplace-knowledge/ [Confirmed at source — (a)]
- Frontiers in Behavioral Economics (2024). Meta-analysis of downsizing and market performance, encompassing 905 effect sizes. [Peer-reviewed; confirmed via secondary reporting — (b)]